

Resource Buying Games

Tobias Harks*

Britta Peis*[†]

Abstract

In *resource buying games* a set of players jointly buys a subset of a finite resource set E (e.g., machines, edges, or nodes in a digraph). The cost of a resource e depends on the number (or load) of players using e , and has to be paid completely by the players before it becomes available. Each player i needs at least one set of a predefined family $\mathcal{S}_i \subseteq 2^E$ to be available. Thus, resource buying games can be seen as a variant of congestion games in which the load-dependent costs of the resources can be shared arbitrarily among the players. A strategy of player i in resource buying games is a tuple consisting of one of i 's desired configurations $S_i \in \mathcal{S}_i$ together with a payment vector $p_i \in \mathbb{R}_+^E$ indicating how much i is willing to contribute towards the purchase of the chosen resources. In this paper, we study the existence and computational complexity of pure Nash equilibria (PNE, for short) of resource buying games. In contrast to classical congestion games for which equilibria are guaranteed to exist, the existence of equilibria in resource buying games strongly depends on the underlying structure of the families \mathcal{S}_i and the behavior of the cost functions. We show that for marginally non-increasing cost functions, matroids are exactly the right structure to consider, and that resource buying games with marginally non-decreasing cost functions always admit a PNE.

1 Introduction

We introduce and study *resource buying games* as a means to model selfish behavior of players jointly designing a resource infrastructure. In a resource buying game, we are given a finite set N of players and a finite set of resources E . We do not specify the type of the resources, they can be just anything (e.g., edges or nodes in a digraph, processors, trucks, etc.). In our model, the players jointly buy a subset of the resources. Each player $i \in N$ has a predefined family of subsets (called *configurations*) $\mathcal{S}_i \subseteq 2^E$ from which player i needs at least one set $S_i \in \mathcal{S}_i$ to be available. For example, the families \mathcal{S}_i could be the collection of all paths linking two player-specific terminal-nodes s_i, t_i in a digraph $G = (V, E)$, or \mathcal{S}_i could stand for the set of machines on which i can process her job on. The cost c_e of a resource $e \in E$ depends on the number of players using e , and needs to be paid completely by the players before it becomes available. As usual, we assume that the cost functions c_e are non-decreasing and normalized in the sense that c_e never decreases with increasing load, and that c_e is zero if none of the players is using e . In a weighted variant of resource buying games, each player has a specific weight (demand) d_i , and the cost c_e depends on the sum of demands of players using e . In resource buying games, a strategy of player i can be regarded as a tuple (S_i, p_i) consisting of one of i 's desired sets $S_i \in \mathcal{S}_i$, together with a payment vector $p_i \in \mathbb{R}_+^E$ indicating how much i is willing to contribute towards the purchase of the resources. The goal of each player is to pay as little as possible by ensuring that the bought resources contain at least one of her desired configurations. A *pure strategy Nash equilibrium* (PNE, for short) is a strategy profile

*Maastricht University, Email: t.harks@maastrichtuniversity.nl.

[†]Technical University Berlin, Email: peis@math.tu-berlin.de.

$\{(S_i, p_i)\}_{i \in N}$ such that none of the players has an incentive to switch her strategy given that the remaining players stick to the chosen strategy. A formal definition of the model will be given in Section 2.

Previous Work.

As the first seminal paper in the area of resource buying games, Anshelevich et al. [5] introduced *connection games* to model selfish behavior of players jointly designing a network infrastructure. In their model, one is given an undirected graph $G = (V, E)$ with non-negative (fixed) edge costs $c_e, e \in E$, and the players jointly design the network infrastructure by buying a subgraph $H \subseteq G$. An edge e of E is bought if the payments of the players for this edge cover the cost c_e , and, a subgraph H is bought if every $e \in H$ is bought. Each player $i \in N$ has a specified source node $s_i \in V$ and terminal node $t_i \in V$ that she wants to be connected in the bought subgraph. A strategy of a player is a payment vector indicating how much she contributes towards the purchase of each edge in E . Anshelevich et al. show that these games have a PNE if all players connect to a common source. They also show that general connection games might fail to have a PNE (see also Section 1 below). Several follow-up papers (cf. [3, 4, 6, 7, 9, 11, 12]) study the existence and efficiency of pure Nash and strong equilibria in connection games and extensions of them. In contrast to these works, our model is more general as we assume load-dependent congestion costs and weighted players. Load-dependent cost functions play an important role in many real-world applications as, in contrast to fixed cost functions, they take into account the intrinsic coupling between the quality or cost of the resources and the resulting demand for it. A prominent example of this coupling arises in the design of telecommunication networks, where the installation cost depends on the installed bandwidth which in turn should match the demand for it.

Hoefler [13] studied resource buying games for load-dependent non-increasing marginal cost functions generalizing fixed costs. He considers unweighted congestion games modeling cover and facility location problems. Among other results regarding approximate PNEs and the price of anarchy/stability, he gives a polynomial time algorithm computing a PNE for the special case, where every player wants to cover a single element.

First Insights.

Before we describe our results and main ideas in detail, we give two examples motivating our research agenda.




Fig. 1(a)




Fig. 1 (b)

Consider the scheduling game illustrated in Fig.1(a) with two resources (machines) $\{e, f\}$ and three players $\{1, 2, 3\}$ each having unit-sized jobs. Any job fits on any machine, and the processing cost of machines e, f is given by $c_j(\ell_j(S))$, where $\ell_j(S)$ denotes the number of jobs on machine $j \in \{e, f\}$ under schedule S . In our model, each player chooses a strategy which is a tuple consisting of one of the two machines, together with a payment vector indicating how much she is willing to pay for each of the machines. Now, suppose the cost functions

for the two machines are $c_e(0) = c_f(0) = 0$, $c_e(1) = c_f(1) = 1$, $c_e(2) = c_f(2) = 1$ and $c_e(3) = c_f(3) = M$ for some large $M > 0$. One can easily verify that there is no PNE: If two players share the cost of one machine, then a player with positive payments deviates to the other machine. By the choice of M , the case that all players share a single machine can never be a PNE. In light of this quite basic example, we have to restrict the set of feasible cost functions. Although the cost functions c_e and c_f of the machines in this scheduling game are monotonically non-decreasing, their marginal cost function is neither non-increasing, nor non-decreasing, where we call cost function $c_e : \mathbb{N} \rightarrow \mathbb{R}_+$ *marginally non-increasing* [*non-decreasing*] if

$$c_e(x + \delta) - c_e(x) \geq [\leq] c_e(y + \delta) - c_e(y) \quad \forall x \leq y; x, y, \delta \in \mathbb{N}. \quad (1)$$

Note that cost functions with non-increasing marginal costs model economies of scale and include fixed costs as a special case. Now suppose that marginal cost functions are non-increasing and consider scheduling games on restricted machines with uniform jobs. It is not hard to establish a simple polynomial time algorithm to compute a PNE for this setting: Sort the machines with respect to the costs evaluated at load one. Iteratively, let the player whose minimal cost among her available resources is maximal exclusively pay for that resource, drop this player from the list and update the cost on the bought resource with respect to a unit increment of load.

While the above algorithm might give hope for obtaining a more general existence and computability result for PNEs for non-increasing marginal cost functions, we recall a counter-example given by [5]. Consider the connection game illustrated in Fig.1(b), where there are two players that want to establish an s_i - t_i path for $i = 1, 2$. Any strategy profile (*state*) of the game contains two paths, one for each player, that have exactly one edge e in common. In a PNE, no player would ever pay a positive amount for an edge that is not on her chosen path. Now, a player paying a positive amount for e (and at least one such player exists) would have an incentive to switch strategies as she could use the edge that is exclusively used (and paid) by the other player for free. Note that this example uses fixed costs which are marginally non-increasing.

Our Results and Outline.

We study unweighted and weighted resource buying games and investigate the existence and computability of pure-strategy Nash equilibria (PNEs, for short). In light of the examples illustrated in Fig.1, we find that equilibrium existence is strongly related to two key properties of the game: the monotonicity of the marginal cost functions and the combinatorial structure of the allowed strategy spaces of the players.

We first consider non-increasing marginal cost functions and investigate the combinatorial structure of the strategy spaces of the players for which PNEs exist. As our main result we show that *matroids* are exactly the right structure to consider in this setting: In Section 3, we present a polynomial-time algorithm to compute a PNE for unweighted matroid resource buying games. This algorithm can be regarded as a far reaching, but highly non-trivial extension of the simple algorithm for scheduling games described before: starting with the collection of matroids, our algorithm iteratively makes use of deletion and contraction operations to minor the matroids, until a basis together with a suitable payment vector for each of the players is found. The algorithm works not only for fixed costs, but also for the more general marginally non-increasing cost functions. Matroids have a rich combinatorial structure and include, for instance, the setting where each player wants to build a spanning tree in a graph. In Section 4, we study weighted resource buying games. We prove that for

non-increasing marginal costs and matroid structure, every (socially) optimal configuration profile can be obtained as a PNE. The proof relies on a complete characterization of configuration profiles that can appear as a PNE. We lose, however, polynomial running time as computing an optimal configuration profile is NP-hard even for simple matroid games with uniform players. In Section 5, we show that our existence result is "tight" by proving that the matroid property is also the maximal property of the configurations of the players that leads to the existence of a PNE: For every two-player weighted resource buying game having non-matroid set systems, we construct an isomorphic game that does not admit a PNE.

We finally turn in Section 6 to resource buying games having non-decreasing marginal costs. We show that every such game possesses a PNE regardless of the strategy space. We prove this result by showing that an optimal configuration profile can be obtained as a PNE. We further show that one can compute a PNE efficiently whenever one can compute a best response efficiently. Thus, PNE can be efficiently computed even in multi-commodity network games.

Connection to Classical Congestion Games.

We briefly discuss connections and differences between resource buying games and classical congestion games. Recall the congestion game model: the strategy space of each player $i \in N$ consists of a family $\mathcal{S}_i \subseteq 2^E$ of a finite set of resources E . The cost c_e of each resource $e \in E$ depends on the number of players using e . In a classical congestion game, each player i chooses one set $S_i \in \mathcal{S}_i$ and needs to pay the *average* cost of every resource in S_i . Rosenthal [14] proved that congestion games always have a PNE. This stands in sharp contrast to resource buying games for which PNE need not exist even for unweighted singleton two-player games with non-decreasing costs, see Fig.1(a). For congestion games with weighted players, Ackermann et al. [1] showed that for non-decreasing marginal cost functions matroids are the maximal combinatorial structure of strategy spaces admitting PNE. In contrast, Theorem 6.1 shows that resource buying games with non-decreasing marginal cost functions always have a PNE *regardless* of the strategy space. Our characterization of matroids as the maximal combinatorial structure admitting PNE for resource buying games with non-increasing marginal costs is also different to the one of Ackermann et al. [1] for classical weighted matroid congestion games with non-decreasing marginal costs. Ackermann et al. prove the existence of PNE by using a potential function approach. Our existence result relies on a complete characterization of PNE implying that there exist payments so that the optimal profile becomes a PNE. For unweighted matroid congestion games, Ackermann et al. [2] prove polynomial convergence of best-response by using a (non-trivial) potential function argument. Our algorithm and its proof of correctness are completely different relying on matroid minors and cuts.

These structural differences between the two models become even more obvious in light of the computational complexity of computing a PNE. In classical network congestion games with non-decreasing marginal costs it is PLS-hard to compute a PNE [2, 10] even for unweighted players. For network games with weighted players and non-decreasing marginal costs, Dunkel and Schulz [8] showed that it is NP-complete to decide whether a PNE exists. In resource buying (network) games with non-decreasing marginal costs one can compute a PNE in polynomial time even with weighted players (Theorem 6.2).

2 Preliminaries

The Model. A tuple $\mathcal{M} = (N, E, \mathcal{S}, (d_i)_{i \in N}, (c_r)_{r \in E})$ is called a *congestion model*, where $N = \{1, \dots, n\}$ is the set of players, $E = \{1, \dots, m\}$ is the set of *resources*, and $\mathcal{S} = \times_{i \in N} \mathcal{S}_i$ is a set of states (also called *configuration profiles*). For each player $i \in N$, the set \mathcal{S}_i is a non-empty set of subsets $S_i \subseteq E$, called *the configurations of i* . If $d_i = 1$ for all $i \in N$ we obtain an *unweighted* game, otherwise, we have a *weighted* game. We call a configuration profile $S \in \mathcal{S}$ (*socially*) *optimal* if its total cost $c(S) = \sum_{e \in E} c_e(S)$ is minimal among all $S \in \mathcal{S}$.

Given a state $S \in \mathcal{S}$, we define $\ell_e(S) = \sum_{i \in N: e \in S_i} d_i$ as the total load of e in S . Every resource $e \in E$ has a *cost function* $c_e : \mathcal{S} \rightarrow \mathbb{N}$ defined as $c_e(S) = c_e(\ell_e(S))$. In this paper, all cost functions are non-negative, non-decreasing and normalized in the sense that $c_e(0) = 0$. We now obtain a *weighted resource buying game* as the (infinite) strategic game $G = (N, \mathcal{S} \times \mathcal{P}, \pi)$, where $\mathcal{P} = \times_{i \in N} \mathcal{P}_i$ with $\mathcal{P}_i = \mathbb{R}_+^{|E|}$ is the set of feasible payments for the players. Intuitively, each player chooses a configuration $S_i \in \mathcal{S}_i$ and a payment vector p_i for the resources. We say that a resource $e \in E$ is *bought* under strategy profile (S, p) , if $\sum_{i \in N} p_i^e \geq c_e(\ell_e(S))$, where p_i^e denotes the payment of player i for resource e . Similarly, we say that a subset $T \subseteq E$ is bought if every $e \in T$ is bought. The private cost function of each player $i \in N$ is defined as $\pi_i(S) = \sum_{e \in E} p_i^e$ if S_i is bought, and $\pi_i(S) = \infty$, otherwise. We are interested in the existence of pure Nash equilibria, i.e., strategy profiles that are resilient against unilateral deviations. Formally, a strategy profile (S, p) is a *pure Nash equilibrium*, PNE for short, if $\pi_i(S, p) \leq \pi_i((S'_i, S_{-i}), (p'_i, p_{-i}))$ for all players $i \in N$ and all strategies $(S_i, p_i) \in \mathcal{S}_i \times \mathcal{P}_i$. Note that for PNE, we may assume w.l.o.g that a pure strategy (S_i, p_i) of player i satisfies $p_i^e \geq 0$ for all $e \in S_i$ and $p_i^e = 0$, else.

Matroid Games. We call a weighted resource buying game a *matroid (resource buying) game* if each configuration set $\mathcal{S}_i \subseteq 2^{E_i}$ with $E_i \subseteq E$ forms the base set of some matroid $\mathcal{M}_i = (E_i, \mathcal{S}_i)$. As it is usual in matroid theory, we will throughout write \mathcal{B}_i instead of \mathcal{S}_i , and \mathcal{B} instead of \mathcal{S} , when considering matroid games. Recall that a non-empty anti-chain¹ $\mathcal{B}_i \subseteq 2^{E_i}$ is the base set of a matroid $\mathcal{M}_i = (E_i, \mathcal{B}_i)$ on resource (*ground*) set E_i if and only if the following *basis exchange property* is satisfied: whenever $X, Y \in \mathcal{B}_i$ and $x \in X \setminus Y$, then there exists some $y \in Y \setminus X$ such that $X \setminus \{x\} \cup \{y\} \in \mathcal{B}_i$. For more about matroid theory, the reader is referred to [15].

3 An Algorithm for Unweighted Matroid Games

Let $\mathcal{M} = (N, E, \mathcal{B}, (c_e)_{e \in E})$ be a model of an unweighted matroid resource buying game. Thus, $\mathcal{B} = \times_{i \in N} \mathcal{B}_i$ where each \mathcal{B}_i is the base set of some matroid $\mathcal{M}_i = (E_i, \mathcal{B}_i)$, and $E = \bigcup_{i \in N} E_i$. In this section, we assume that the cost functions c_e , $e \in E$ are marginally non-increasing.

Given a matroid $\mathcal{M}_i = (E_i, \mathcal{B}_i)$, we denote by $\mathcal{I}_i = \{I \subseteq E \mid I \subseteq B \text{ for some } B \in \mathcal{B}_i\}$ the collection of *independent sets* in \mathcal{M}_i . Furthermore, we call a set $C \subseteq E_i$ a *cut* of matroid \mathcal{M}_i if $E_i \setminus C$ does not contain a basis of \mathcal{M}_i . Let $\mathcal{C}_i(\mathcal{M}_i)$ denote the collection of all inclusion-wise minimal cuts of \mathcal{M}_i . We will need the following basic insight at several places.

Lemma 3.1. [15, Chapters 39 – 42] *Let \mathcal{M} be a weighted matroid with weight function $w : E \rightarrow \mathbb{R}_+$. A basis B is a minimum weight basis of \mathcal{M} if and only if there exists no basis B^* with $|B \setminus B^*| = 1$ and $w(B^*) < w(B)$.*

¹Recall that $\mathcal{B}_i \subseteq 2^{E_i}$ is an *anti-chain* (w.r.t. $(2^{E_i}, \subseteq)$) if $B, B' \in \mathcal{B}_i$, $B \subseteq B'$ implies $B = B'$.

In a strategy profile (B, p) of our game with $B = (B_1, \dots, B_n) \in \mathcal{B}$ (and $n = |N|$) players will jointly buy a subset of resources $\bar{B} \subseteq E$ with $\bar{B} = B_1 \cup \dots \cup B_n$. Such a strategy profile (B, p) is a PNE if none of the players $i \in N$ would need to pay less by switching to some other basis $B'_i \in \mathcal{B}_i$, given that all other players $j \neq i$ stick to their chosen strategy (B_j, p_j) . By Lemma 3.1, it suffices to consider bases $\hat{B}_i \in \mathcal{B}_i$ with $\hat{B}_i = B_i - g + f$ for some $g \in B_i \setminus \hat{B}_i$ and $f \in \hat{B}_i \setminus B_i$. Note that by switching from B_i to \hat{B}_i , player i would need to pay the additional marginal cost $c_f(l_f(B) + 1) - c_f(l_f(B))$, but would not need to pay for element g . Thus, (B, p) is a PNE iff for all $i \in N$ and all $\hat{B}_i \in \mathcal{B}_i$ with $\hat{B}_i = B_i - g + f$ for some $g \in B_i \setminus \hat{B}_i$ and $f \in \hat{B}_i \setminus B_i$ holds $p_i^g \leq c_f(l_f(B) + 1) - c_f(l_f(B))$.

We now give a polynomial time algorithm (see Algorithm 1 below) computing a PNE for unweighted matroid games with marginally non-increasing costs. The idea of the algorithm can roughly be described as follows: In each iteration, for each player $i \in N$, the algorithm maintains some independent set $B_i \in \mathcal{I}_i$, starting with $B_i = \emptyset$, as well as some payment vector $p_i \in \mathbb{R}_+^E$, starting with the all-zero vector. It also maintains a current matroid $\mathcal{M}'_i = (E'_i, \mathcal{B}'_i)$ that is obtained from the original matroid $\mathcal{M}_i = (E_i, \mathcal{B}_i)$ by deletion and contraction operations (see e.g., [15] for the definition of deletion and contraction in matroids.) The algorithm also keeps track of the current marginal cost $c'_e = c_e(\ell_e(B) + 1) - c_e(\ell_e(B))$ for each element $e \in E$ and the current sequence $B = (B_1, \dots, B_n)$. Note that c'_e denotes the amount that needs to be paid if some additional player i selects e into its set B_i . In each iteration, while there exists at least one player i such that B_i is not already a basis, the algorithm chooses among all cuts in $\mathcal{C} = \{C \in \mathcal{C}_i(\mathcal{M}'_i) \mid \text{for some } i \in N\}$ an inclusion-wise minimal cut C^* whose bottleneck element (i.e., the element of minimal current weight in C^*) has maximal c' -weight (step 3). (We assume that some fixed total order (E, \preceq) is given to break ties, so that the choices of C^* and e^* are unique.) It then selects the bottleneck element $e^* \in C^*$ (step 4), and some player i^* with $C^* \in \mathcal{C}_i(\mathcal{M}'_{i^*})$ (step 5). In an update step, the algorithm lets player i^* pay the marginal cost c'_{e^*} (step 7), adds e^* to B_{i^*} (step 8), and contracts element e^* in matroid \mathcal{M}'_{i^*} (step 12). If B_{i^*} is a basis in the original matroid \mathcal{M}_{i^*} , the algorithm drops i^* from the player set N (step 10). Finally, the algorithm deletes the elements in $C^* \setminus \{e^*\}$ in all matroids \mathcal{M}'_i for $i \in N$ (step 16), and iterates until $N = \emptyset$, i.e., until a basis has been found for all players.

Obviously, the algorithm terminates after at most $|N| \cdot |E|$ iterations, since in each iteration, at least one element e^* is dropped from the ground set of one of the players. Note that the inclusion-wise minimal cut C^* whose bottleneck element e^* has maximal weight (step 3), as well as the corresponding player i^* and the bottleneck element e^* , can be efficiently found, see the appendix for a corresponding subroutine.

It is not hard to see that Algorithm 1 corresponds exactly to the procedure described in Section 1 to solve the scheduling game (i.e., the matroid game on uniform matroids) with non-increasing marginal cost functions. We show that the algorithm returns a pure Nash equilibrium also for general matroids. As a key Lemma, we show that the current weight of the chosen bottleneck element monotonically decreases.

Theorem 3.1. *The output (B, p) of the algorithm is a PNE.*

Proof. Obviously, at termination, each set B_i is a basis of matroid \mathcal{M}_i , as otherwise, player i would not have been dropped from N , in contradiction to the stopping criterium $N = \emptyset$. Thus, we first need to convince ourselves that the algorithm terminates, i.e., constructs a basis B_i for each matroid \mathcal{M}_i . However, this follows by the definition of contraction and deletion in matroids:

Algorithm 1 COMPUTING PNE IN MATROIDS

Input: $(N, E, \mathcal{M}_i = (E_i, \mathcal{B}_i), c)$ **Output:** PNE (B, p)

```
1: Initialize  $\mathcal{B}'_i = \mathcal{B}_i, E'_i = E_i, B_i = \emptyset, p_i^e = 0, t_e = 1$ , and  $c'_e = c_e(1)$  for each  $i \in N$  and each  $e \in E$ ;  
2: while  $N \neq \emptyset$  do  
3:   choose  $C^* \leftarrow \operatorname{argmax}\{\min\{c'_e : e \in C\} \mid C \in \mathcal{C} \text{ inclusion-wise minimal}\}$   
   where  $\mathcal{C} = \{C \in \mathcal{C}_i(\mathcal{M}'_i) \mid \text{for some player } i \in N\}$ ;  
4:   choose  $e^* \leftarrow \operatorname{argmin}\{c'_e \mid e \in C^*\}$ ;  
5:   choose  $i^*$  with  $C^* \in \mathcal{C}_{i^*}(\mathcal{M}'_{i^*})$ ;  
6:    $p_{i^*}^{e^*} \leftarrow c'_{e^*}$ ;  
7:    $c'_{e^*} \leftarrow c_{e^*}(t_{e^*} + 1) - c_{e^*}(t_{e^*})$ ;  
8:    $B_{i^*} \leftarrow B_{i^*} + e^*$ ;  
9:   if  $B_{i^*} \in \mathcal{B}_{i^*}$  then  
10:     $N \leftarrow N - i^*$ ;  
11:   end if  
12:    $\mathcal{B}'_{i^*} \leftarrow \mathcal{B}'_{i^*}/e^* = \{B \subseteq E'_{i^*} \setminus \{e^*\} \mid B + e^* \in \mathcal{B}'_{i^*}\}$ ;  
13:    $E'_{i^*} \leftarrow E'_{i^*} \setminus \{e^*\}$ ;  
14:    $t_{e^*} \leftarrow t_{e^*} + 1$ ;  
15:   for all players  $i \in N$  do  
16:      $\mathcal{B}'_i \leftarrow \mathcal{B}'_i \setminus (C^* \setminus \{e^*\}) = \{B \subseteq E'_i \setminus (C^* \setminus \{e^*\}) \mid B \in \mathcal{B}'_i\}$   
17:      $E'_i \leftarrow E'_i \setminus (C^* \setminus \{e^*\})$ ;  
18:   end for  
19: end while  
20:  $B = (B_1, \dots, B_n), p = (p_1, \dots, p_n)$ ;  
21: Return  $(B, p)$ 
```

To see this, we denote by $N^{(k)}$ the current player set, and by $B_i^{(k)}$ and $\mathcal{M}_i^{(k)} = (E_i^{(k)}, \mathcal{B}_i^{(k)})$ the current independent set and matroid of player i at the beginning of iteration k . Suppose that the algorithm now chooses e^* in step 4 and player i^* in step 5. Thus, it updates $B_{i^*}^{(k+1)} \leftarrow B_{i^*}^{(k)} + e^*$ in step 8 and considers the base set $\mathcal{B}_{i^*}^{(k)}/e^*$ of the contracted matroid $\mathcal{M}_{i^*}^{(k)}/e^*$. Note that for each $B \in \mathcal{B}_{i^*}^{(k)}/e^*$, the set $B + e^*$ is a basis in $\mathcal{B}_{i^*}^{(k)}$, and, by induction, $B + B_{i^*}^{(k+1)}$ is a basis in the original matroid \mathcal{M}_{i^*} . Thus, $B_{i^*}^{(k+1)}$ is a basis in \mathcal{M}_{i^*} (and i^* is dropped from $N^{(k)}$) if and only if $\mathcal{B}_{i^*}^{(k)}/e^* = \{\emptyset\}$.

Now consider any other player $i \neq i^*$ with $\mathcal{B}_i^{(k)} \neq \{\emptyset\}$ (and thus $i \in N^{(k)}$). Then, for the new base set $\mathcal{B}_i^{(k+1)} = \mathcal{B}_i^{(k)} \setminus (C^* \setminus \{e^*\})$ we still have $\mathcal{B}_i^{(k+1)} \neq \{\emptyset\}$, since otherwise $C^* \setminus \{e^*\}$ is a cut in matroid $\mathcal{M}_i^{(k)}$, in contradiction to the choice of C^* . Thus, since the algorithm only terminates when $N^{(k)} = \emptyset$ for the current iteration k , it terminates with a basis B_i for each player i .

Note that throughout the algorithm it is guaranteed that the current payment vectors $p = (p_1, \dots, p_n)$ satisfy $\sum_{i \in N} p_i^e = c_e(\ell_e(B))$ for each $e \in E$ and the current independent sets $B = (B_1, \dots, B_n)$. This follows, since the payments are only modified in step 7, where the marginal payment $p_{i^*}^{e^*} = c_{e^*}(\ell_{e^*}(B) + 1) - c_{e^*}(\ell_{e^*}(B))$ is assigned just before e^* was selected into the set B_{i^*} . Since we assumed the c_e 's to be non-decreasing, this also guarantees that each component p_i^e is non-negative, and positive only if $e \in B_i$.

It remains to show that the final output (B, p) is a PNE. Suppose, for the sake of contradiction, that this were not true, i.e., that there exists some $i \in N$ and some basis $\hat{B}_i \in \mathcal{B}_i$ with

$\hat{B}_i = B_i - g + f$ for some $g \in B_i \setminus \hat{B}_i$ and $f \in \hat{B}_i \setminus B_i$ such that $p_i^g > c_f(l_f(B+1)) - c_f(l_f(B))$. Let k be the iteration in which the algorithm selects the element g to be paid by player i , i.e., the algorithm updates $B_i^{(k+1)} \leftarrow B_i^{(k)} + g$. Let $C^* = C(k)$ be the cut for matroid $\mathcal{M}_i^{(k)} = (E_i^{(k)}, \mathcal{B}_i^{(k)})$ chosen in this iteration. Thus, the set $E_i^{(k)} \setminus C^*$ contains no basis in $\mathcal{B}_i^{(k)}$, i.e., no set $B \subseteq E_i^{(k)} \setminus C^*$ with $B + B_i^{(k)} \in \mathcal{B}_i$. Note that the final set B_i contains no element from C^* other than g , as all elements in $C^* \setminus \{g\}$ are deleted from matroid $\mathcal{M}_i^{(k)}/g$. We distinguish the two cases where $f \in C^*$, and where $f \notin C^*$.

In the first case, if $f \in C^*$, then, since the algorithm chooses g of minimal current marginal weight, we know that $p_i^g = c_g(l_g(B^{(k)} + 1) - c_g(l_g(B^{(k)}))) \leq c_f(l_f(B^{(k)} + 1) - c_f(l_f(B^{(k)})))$. Thus, the marginal cost of f must decrease at some later point in time, i.e., $c_f(l_f(B+1)) - c_f(l_f(B)) < c_f(l_f(B^{(k)} + 1) - c_f(l_f(B^{(k)})))$. But this cannot happen, since f is deleted from all matroids for which the algorithm has not found a basis up to iteration k .

However, also the latter case cannot be true: Suppose $f \notin C^*$. If $f \in E_i^{(k)}$, then $\hat{B}_i \setminus B_i^{(k)} \subseteq E_i^{(k)} \setminus C^*$, but $\hat{B}_i = \hat{B}_i \setminus B_i^{(k)} + B_i^{(k)} \in \mathcal{B}_i$, in contradiction to C^* being a cut in $\mathcal{M}_i^{(k)}$. Thus, f must have been dropped from E_i in some iteration l prior to k by either some deletion or contraction operation. We show that this is impossible (which finishes the proof): A contraction operation of type $\mathcal{M}_i^{(l)} \rightarrow \mathcal{M}_i^{(l)}/e_l$ drops only the contracted element e_l from player i 's ground set $E_i^{(l)}$, after e_l has been added to the current set $B_i^{(l)} \subseteq B_i$. Thus, since $f \notin B_i$, f must have been dropped by the deletion operation in iteration l . Let $C(l)$ be the chosen cut in iteration l , and e_l the bottleneck element. Thus, $f \in C(l) - e_l$. Now, consider again the cut $C^* = C(k)$ of player i which was chosen in iteration k . Recall that the bottleneck element of $C(k)$ in iteration k was g . Note that there exists some cut $C' \supseteq C(k)$ such that C' is a cut of player i in iteration l and $C(k)$ was obtained from C' by the deletion and contraction operations in between iterations l and k . Why did the algorithm choose $C(l)$ instead of C' ? The only possible answer is, that the bottleneck element a of C' has current weight $c_a^{(l)} \leq c_{e_l}^{(l)} \leq c_f^{(l)}$. On the other hand, if f was dropped in iteration l , then $c_f^{(l)} = c_f(l_f(B+1)) - c_f(l_f(B))$. Thus, by our assumption, $c_f^{(l)} < p_i^g = c_g^{(k)}$. However, since the cost function c_g is the marginally non-increasing, it follows that $c_g^{(k)} \leq c_g^{(l)}$. Summarizing, we yield $c_a^{(l)} \leq c_{e_l}^{(l)} \leq c_f^{(l)} < c_g^{(k)} \leq c_g^{(l)}$, and, in particular, $c_{e_l}^{(l)} < c_g^{(k)}$, in contradiction to Lemma 3.2 below (proven in the appendix). \square

Lemma 3.2. *Let \hat{c}_k denote the current weight of the bottleneck element chosen in step 4 of iteration k . Then this weight monotonically decreases, i.e., $l < k$ implies $\hat{c}_l \geq \hat{c}_k$ for all $l, k \in \mathbb{N}$.*

4 Weighted Matroid Games

For proving the existence of PNE in *weighted* matroid games with non-increasing marginal costs our algorithm presented before does not work anymore. We prove, however, that there exists a PNE in matroid games with non-increasing marginal costs even for weighted demands. To obtain our existence result, we now derive a complete characterization of configuration profiles $B \in \mathcal{B}$ in weighted matroid games $(N, E, \mathcal{B}, d, c)$ that can be obtained as a PNE. For our characterization, we need a few definitions: For $B \in \mathcal{B}$, $e \in E$ and $i \in N_e(B) := \{i \in N \mid e \in B_i\}$ let $\text{ex}_i^e := \{f \in E - e \mid B_i - e + f \in \mathcal{B}_i\} \subseteq E$ denote the set of all resources f such that player i could exchange the resources e and f to obtain an alternative basis $B_i - e + f \in \mathcal{B}_i$. Note that $\text{ex}_i(e)$ might be empty, and that, if $\text{ex}_i(e)$ is empty, the element e lies in every basis of player i (by the matroid basis exchange property). Let

$F := \{e \in E \mid e \text{ lies in each basis of } i \text{ for some } i \in N\}$ denote the set of elements that are “fixed” in the sense that they must lie in one of the players’ chosen basis. Furthermore, we define for all $e \in E - F$ and all $i \in N_e(B)$ and all $f \in \text{ex}_i(e)$ the value $\Delta_i(B; e \rightarrow f) := c_f(\ell_f(B_i + f - e, B_{-i})) - c_f(\ell_f(B))$ which is the marginal amount that needs to be paid in order to buy resource f if i switches from B_i to $B_i - e + f$. Finally, let $\Delta_i^e(B)$ be the minimal value among all $\Delta_i(B; e \rightarrow f)$ with $f \in \text{ex}_i(e)$. The proof of the following characterization can be found in the appendix.

Theorem 4.1. *Consider a weighted matroid resource buying game $(N, E, \mathcal{B}, d, c)$. There is a payment vector p such that the strategy profile (B, p) with $B \in \mathcal{B}$ is a PNE if and only if*

$$c_e(B) \leq \sum_{i \in N_e(B)} \Delta_i^e(B) \quad \text{for all } e \in E \setminus F. \quad (2)$$

Note that the above characterization holds for arbitrary non-negative and non-decreasing cost functions. In particular, if property (2) were true, it follows from the constructive proof that the payment vector p can be efficiently computed. The following Theorem 4.2 states that matroid games with non-increasing marginal costs and weighted demands always possess a PNE. We prove Theorem 4.2 (in the appendix) by showing that any socially optimal configuration $B \in \mathcal{B}$ satisfies (2).

Theorem 4.2. *Every weighted matroid resource buying game with marginally non-increasing cost functions possesses a PNE.*

Note that the above existence result does not imply an efficient algorithm for computing a PNE: It is straightforward to show that computing a socially optimal configuration profile is NP-hard even for unit demands and singleton strategies.

5 Non-Matroid Strategy Spaces

In the previous section, we proved that for weighted matroid congestion games with non-negative, non-decreasing, marginally non-increasing cost functions, there always exists a PNE. In this section, we show that the matroid property of the configuration sets is also the maximal property needed to guarantee the existence of a PNE for *all* weighted resource buying games with marginally non-increasing costs (assuming that there is no a priori combinatorial structure how the strategy spaces are interweaved). This result and its proof (in the appendix) is related to one of Ackermann et al. in [1] for the classical weighted matroid congestion games with average cost sharing and marginally non-decreasing cost functions.

Theorem 5.1. *For every non-matroid anti-chain \mathcal{S} on a set of resources E , there exists a weighted two-player resource buying game $G = (\tilde{E}, (\mathcal{S}_1 \times \mathcal{S}_2) \times \mathcal{P}, \pi)$ having marginally non-increasing cost functions, whose strategy spaces \mathcal{S}_1 and \mathcal{S}_2 are both isomorphic to \mathcal{S} , so that G does not possess a PNE.*

6 Non-Decreasing Marginal Cost Functions

In this section, we consider non-decreasing marginal cost functions on weighted resource buying games in general, i.e., $\mathcal{S} = \times_{i \in N} \mathcal{S}_i$ is not necessarily the cartesian product of matroid base sets anymore. We prove that for every socially optimal state S^* in a congestion model with non-decreasing marginal costs, we can define *marginal cost* payments p^* that result in

a PNE. Formally, for a given socially optimal configuration profile $S^* \in \mathcal{S}$ and a fixed order $\sigma = 1, \dots, n$ of the players, we let $N_e(S^*) := \{i \in N \mid e \in S_i^*\}$ denote the players using e in S^* , $N_e^j(S^*) := \{i \in N_e(S^*) \mid i \leq_\sigma j\}$ denote the players in $N_e(S^*)$ prior or equal to j in σ , and $\ell_e^{\leq j}(S^*) = \sum_{i \in N_e^j(S^*)} d_i$ denote the load of these players on e in S^* . Given these definitions, we allocate the cost $c_e(\ell_e(S^*))$ for each resource $e \in E$ among the players in $N_e(S^*)$ by setting $p_i^e = 0$ if $e \notin S_i^*$ and $p_i^e = c_e(\ell_e^{\leq j}(S^*)) - c_e(\ell_e^{\leq j-1}(S^*))$ if player i is the j -th player in $N_e(S^*)$ w.r.t. σ . Let us call this payment vector *marginal cost pricing*.

Theorem 6.1. *Let S^* be a socially optimal solution. Then, marginal cost pricing induces a PNE. (The proof is in the appendix.)*

We now show that there is a simple polynomial time algorithm computing a PNE whenever we are able to efficiently compute a best-response. By simply inserting the players one after the other using their current best-response with respect to the previously inserted players, we obtain a PNE. It follows that for (multi-commodity) network games we can compute a PNE in polynomial time.

Theorem 6.2. *For multi-commodity network games with non-decreasing marginal costs, there is a polynomial time algorithm computing a PNE.*

The proof is straight-forward: Because payments of previously inserted players do not change in later iterations and marginal cost functions are non-decreasing, the costs of alternative strategies only increase as more players are inserted. Thus, the resulting strategy profile is a PNE.

7 Conclusions and Open Questions

We presented a detailed study on the existence and computational complexity of pure Nash equilibria in resource buying games. Our results imply that the price of stability is always one for both, games with non-decreasing marginal costs and games with non-increasing marginal costs and matroid structure. Regarding the price of anarchy, even on games with matroid structure it makes a difference whether cost functions are marginally non-increasing or marginally non-decreasing. For non-increasing marginal costs it is known that the price of anarchy is n (the lower bound even holds for singleton games). On the other hand, for non-decreasing marginal costs we can show that the price of anarchy for uniform and partition matroids is exactly n , while it is unbounded for graphical matroids even on instances with only two players. Convergence of best-response dynamics has not been addressed so far and deserves further research.

References

- [1] H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. *Theor. Comput. Sci.*, 410(17):1552–1563, 2009.
- [2] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55:25:1–25:22, December 2008.
- [3] Elliot Anshelevich and Bugra Caskurlu. Exact and approximate equilibria for optimal group network formation. In *ESA*, pages 239–250, 2009.

- [4] Elliot Anshelevich and Bugra Caskurlu. Price of stability in survivable network design. In *Proc. of SAGT*, pages 208–219, 2009.
- [5] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- [6] Elliot Anshelevich and Adriana Karagiozova. Terminal backup, 3d matching, and covering cubic graphs. *SIAM J. Comput.*, 40(3):678–708, 2011.
- [7] Jean Cardinal and Martin Hoefer. Non-cooperative facility location and covering games. *Theor. Comput. Sci.*, 411:1855–1876, March 2010.
- [8] J. Dunkel and A. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Math. Oper. Res.*, 33(4):851–868, 2008.
- [9] Amir Epstein, Michal Feldman, and Yishay Mansour. Strong equilibrium in cost sharing connection games. *Games Econom. Behav.*, 67(1):51–68, September 2009.
- [10] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. of the 36th STOC*, pages 604–612, 2004.
- [11] Martin Hoefer. Non-cooperative tree creation. *Algorithmica*, 53:104–131, 2009.
- [12] Martin Hoefer. Strategic cooperation in cost sharing games. In *Proc. of WINE*, pages 258–269, 2010.
- [13] Martin Hoefer. Competitive cost sharing with economies of scale. *Algorithmica*, 60:743–765, 2011.
- [14] R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory*, 2(1):65–67, 1973.
- [15] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003.

A Omitted Proofs

A.1 A Subroutine to Detect C^*, e^* and i^*

We might assume that the elements of $E' = E'_1 \cup \dots \cup E'_n = \{e_1, \dots, e_m\}$ are listed by non-increasing current c' -weights and such that the tie-breaking rule induced by the total order (E, \preceq) is respected. (Note that in each iteration, only the c' -weight of the chosen bottleneck element might change to some smaller weight. The c' -weight of the remaining elements keeps the same.) The following procedure is used to detect an inclusion-wise minimal cut C^* of some player i^* with the property that the bottleneck element e^* of C^* has maximal c' -weight among all possible cuts in this game: Initially, set $C^* = \{e_m\}$ and $k = m$. If C^* is a cut (i.e., if $E' - C^*$ does not contain a basis) for at least some player, set $e^* = e_k$ and decrease C^* to some inclusion-wise minimal cut as follows: as long as possible, choose some element $e \in C^* - e^*$ such that $C^* - e$ remains a cut for at least one of the players $i^* \in N$, and set $C^* \leftarrow C^* - e$; Return (C^*, e^*, i^*) ; Else, update $C^* \leftarrow C^* + e_{k-1}$, and iterate with $k \leftarrow k - 1$.

A.2 Proof of Lemma 3.2

Proof. For each iteration k let $c_e^{(k)}$ denote the current weight of element e , and $\mathcal{C}_i^{(k)}$ denote the current set of inclusion-wise minimal cuts for player i . Note that for each cut $C \in \mathcal{C}_i^{(k)}$ with $k > 1$, there exists some $C' \in \mathcal{C}_i^{(k-1)}$ such that $C \subseteq C'$ and C is obtained from C' by the contraction and deletion operations of iteration $k - 1$. For the sake of contradiction, suppose that k is the first iteration such that $\hat{c}_k > \hat{c}_{k-1}$. Let e be the bottleneck element chosen in step 4 of iteration $k - 1$. Thus, the corresponding cut $C(k)$ that was chosen in step 3 of iteration k must be obtained from some larger cut C' by removing at least one element $a \in C'$ with $c_a^{(k-1)} \leq \hat{c}_{k-1} = c_e^{(k-1)}$, and, if equality, with $a \prec e$. Since the deletion operation of iteration $k - 1$ removes only elements $e' \in E$ of weight $c_{e'}^{(k-1)} \geq c_e^{(k-1)}$, and if equality, those with $e' \succ e$, the element a must have been dropped from C' by contracting e , i.e., $a = e$. Since this contraction operation touches only the matroid of the player chosen in iteration $k - 1$, say i , it suffices to consider only the cut sets $\mathcal{C}_i^{(k)}$ and $\mathcal{C}_i^{(k-1)}$ and the base sets $\mathcal{B}_i^{(k)}$ and $\mathcal{B}_i^{(k-1)}$ of player i in iterations k and $k - 1$. So far, we observed that $a \in C' \cap C(k - 1)$ where C' and $C(k - 1)$ are both cuts in $\mathcal{C}_i^{(k-1)}$, and that the element a vanishes from cut C' by the contraction operation $\mathcal{M}_i^{(k-1)} \rightarrow \mathcal{M}_i^{(k-1)}/a$. Thus, $C' - a$ must be a (not necessarily inclusion-wise minimal) cut in $\mathcal{M}_i^{(k-1)}/a$. However, since C' is an inclusion-wise minimal cut in $\mathcal{M}_i^{(k-1)}$, the set $E_i^{(k-1)} - (C' - a)$ contains some basis $\hat{B} \in \mathcal{B}_i^{(k-1)}$ with $a \in \hat{B}$. Thus, $B := \hat{B} - a$ is a set in $E_i^{(k-1)} - (C' - a)$ with $B + a \in \mathcal{B}_i^{(k-1)}$, in contradiction to $C' - a$ being a cut in $\mathcal{M}_i^{(k-1)}/a$. \square

A.3 Proof of Theorem 4.1

Proof. We first proof the "only if" direction. Let (B, p) be a PNE. Then, by Lemma 3.1 and the definition of a PNE, we obtain for all $e \in E \setminus F$:

$$c_e(B) = \sum_{i \in N_e(B)} p_i^e \leq \sum_{i \in N_e(B)} \Delta_i^e(B).$$

Note that the $\Delta_i^e(B)$ are well defined as we only consider elements in $E \setminus F$. Now we prove the "if" direction. For all $e \in F$ we pick a player i with $\text{ex}_i(e) = \emptyset$ and let her pay the entire cost, i.e., $p_i^e = c_e(B)$. For all $e \in E \setminus F$ and $i \in N_e(B)$, we define

$$p_i^e = \frac{\Delta_i^e(B)}{\sum_{j \in N_e(B)} \Delta_j^e(B)} \cdot c_e(B),$$

if the denominator is positive, and $p_i^e = 0$, otherwise. Using (2), we obtain

$$p_i^e \leq \Delta_i^e(B) \text{ for all } e \in E \setminus F$$

proving that (B, p) is a PNE. \square

A.4 Proof of Theorem 4.2

Proof. We prove that any socially optimal configuration profile $B \in \mathcal{B}$ satisfies (2) and, thus, by Theorem 4.1 there exists a payment vector p such that (B, p) is a PNE. Assume by contradiction that B does not satisfy (2). Hence, there is an $e \in E \setminus F$ with

$$c_e(B) > \sum_{i \in N_e(B)} \Delta_i^e(B). \quad (3)$$

By relabeling indices we may write $N_e(B) = \{1, \dots, k\}$ for some $1 \leq k \leq n$, and define for every $i \in N_e(B)$ the tuple $(\hat{B}_i, f_i) \in \mathcal{B}_i \times (E_i - e)$ as the one minimizing $\Delta_i(B; e \rightarrow f)$ among all tuples $(B'_i, f) \in \mathcal{B}_i \times (E_i - e)$ with $B'_i = B_i + f - e \in \mathcal{B}_i$. Note that (\hat{B}_i, f_i) is well defined as $e \in E \setminus F$. We now iteratively change the current basis of every player in $N_e(B)$ in the order of their indices to the alternative basis $\hat{B}_i, i = 1, \dots, k$. This gives a sequence of profiles (B^0, B^1, \dots, B^k) with $B^0 = B$ and $B^i = (\hat{B}_i, B_{-i}^{i-1})$ for $i = 1, \dots, k$. For the cost increase of the new elements $f_i, i \in N_e(B)$, we obtain the key inequality $c_{f_i}(\ell_{f_i}(B^{i-1})) - c_{f_i}(\ell_{f_i}(B^i)) \leq \Delta_i^e(B)$. This inequality holds because cost functions are marginally non-increasing, that is, the marginal costs only decrease with higher load. Plugging everything together, yields

$$\begin{aligned} c(B) - c(B^k) &= \sum_{i=1}^k (c(B^{i-1}) - c(B^i)) = \sum_{i=1}^k (c_e(\ell_e(B^{i-1})) + c_{f_i}(\ell_{f_i}(B^{i-1})) \\ &\quad - c_e(\ell_e(B^i)) - c_{f_i}(\ell_{f_i}(B^i))) \\ &= c_e(\ell_e(B)) - c_e(\ell_e(B^k)) + \sum_{i=1}^k (c_{f_i}(\ell_{f_i}(B^{i-1})) - c_{f_i}(\ell_{f_i}(B^i))) \\ &\geq c_e(\ell_e(B)) - \sum_{i=1}^k \Delta_i^e(B) > 0, \end{aligned}$$

where the first inequality uses $c_e(\ell_e(B^k)) = c_e(0) = 0$ (note that $e \in E \setminus F$) and the assumption that cost functions have non-increasing marginal costs. The second strict inequality follows from (3). Altogether, we obtain a contradiction to the optimality of B . \square

A.5 Proof of Theorem 5.1

Recall that $\mathcal{S} \subseteq 2^E$ is an *anti-chain* (with respect to $(2^E, \subseteq)$) if for every $X \in \mathcal{S}$, no proper superset $Y \subset X$ belongs to \mathcal{S} . Also note that it suffices to consider configuration sets \mathcal{S}_i that form an anti-chain, as (due to the non-negative cost functions) player i would never have an incentive to switch her strategy to a superset of her chosen one.

We call \mathcal{S} a *non-matroid* set system if the tuple $(E, \{X \subseteq S : S \in \mathcal{S}\})$ is not a matroid. The following Lemma can also be derived from the proof of Lemma 16 in [1].

Lemma A.1. *If $\mathcal{S} \subseteq 2^E$ is a non-matroid antichain, then there exist $X, Y \in \mathcal{S}$ and $\{a, b, c\} \subseteq X \cup Y$ such that each set $Z \subseteq (X \cup Y) - a$ contains both, b and c .*

Proof. Recall the *basis exchange property* for matroids: an anti-chain $\mathcal{B} \subseteq 2^E$ is the family of bases of some matroid if and only if for any $X, Y \in \mathcal{B}$ and $x \in X \setminus Y$ there exists some $y \in Y \setminus X$ such that $X - x + y \in \mathcal{B}$. Thus, if the anti-chain $\mathcal{S} \subseteq 2^E$ is a non-matroid, there must exist $X, Y \in \mathcal{S}$ and $x \in X \setminus Y$ such that for all $y \in Y \setminus X$ the set $X - x + y$ does not belong to \mathcal{S} . We choose such X, Y and $x \in X \setminus Y$ with $|Y \setminus X|$ minimal (among all $Y' \in \mathcal{S}$ with $X - x + y' \notin \mathcal{S}$ for all $y' \in Y' \setminus X$). We distinguish the two cases $|Y \setminus X| = 1$ and $|Y \setminus X| \geq 1$: In case $|Y \setminus X| = 1$, set $\{a\} = Y \setminus X$ and choose any two distinct elements $\{b, c\} \in X \setminus Y$. Note that $|X \setminus Y| \geq 2$ as otherwise, if $X \setminus Y = \{x\}$, then $Y = X - x + a$, in contradiction to our assumption. Now, for any set $Z \subseteq (X \cup Y) - a$, the anti-chain property implies $Z = X$, and therefore $\{b, c\} \subseteq Z$, as desired.

In the latter case $|Y \setminus X| \geq 1$, we choose any two distinct elements $\{b, c\} \in Y \setminus X$ and set $a = x$. Consider any $Z \in \mathcal{S}$ with $Z \subseteq (X \cup Y) - a$ and suppose, for the sake of contradiction, that $\{b, c\} \not\subseteq Z$. Since $Z \setminus X \subseteq Y \setminus X$, there cannot exist some $z \in Z \setminus X$ with $X - a + z \in \mathcal{S}$. However, $|Z \setminus X| < |Y \setminus X|$ in contradiction to our choice of Y . \square

Proof of the Theorem. Let $\mathcal{S}_1 \subseteq 2^{E_1}$ and $\mathcal{S}_2 \subseteq 2^{E_2}$ be the two strategy spaces for player one and player two, respectively, both isomorphic to our given non-matroid anti-chain $\mathcal{S} \subseteq 2^E$. In the following, we describe the game G by defining the demands and costs and describing how the resources and strategy spaces interweave: For each player $i = 1, 2$, choose $X_i, Y_i \in \mathcal{S}_i$ and $\{a_i, b_i, c_i\} \subseteq E_i$ as described in Lemma A.1. In our game G , the two players have only three resources in common, i.e., $\{x, y, z\} = E_1 \cap E_2$. We set $x := a_1 = b_2$, $y := a_2 = b_1$ and $z := c_1 = c_2$. All other resources in $E_i \setminus \{x, y, z\}$ are exclusively used by player i for $i = 1, 2$. We define the (load-dependent) costs $c_e(t)$, $t \in \mathbb{R}_+$ for the resources $e \in \tilde{E} = E_1 \cup E_2$ as follows: all elements in $(X_1 \cup X_2 \cup Y_1 \cup Y_2) \setminus \{x, y, z\}$ have a cost of zero, and all elements in $E_1 \setminus (X_1 \cup Y_1)$ and in $E_2 \setminus (X_2 \cup Y_2)$ have some very large cost M . The costs on $\{x, y, z\}$ are defined as $c_x(t) = t$, $c_y(t) = 5\frac{1}{2}$ and $c_z(t) = 4$. Note that each of these cost functions is non-negative, non-decreasing and marginally non-increasing.

Now, suppose that (Z^*, p^*) with $Z^* = (Z_1^*, Z_2^*) \in \mathcal{S}_1 \times \mathcal{S}_2$ and $p^* = (p_1^*, p_2^*) \in \mathbb{R}_+^{E_1} \times \mathbb{R}_+^{E_2}$ were a PNE for the game as described above with demands $d_1 = 5$ and $d_2 = 4$. Choosing M large enough ensures that $Z_i^* \subseteq X_i \cup Y_i$ for each player $i \in \{1, 2\}$. Moreover, by the choice of X_i and Y_i in the proof of Lemma A.1, there exist $S_1, T_1 \in \mathcal{S}_1$ with $x \in S_1$, $\{y, z\} \cap S_1 = \emptyset$ and $x \notin T_1 \supseteq \{y, z\}$, as well as $S_2, T_2 \in \mathcal{S}_2$ with $y \in S_2$, $\{x, z\} \cap S_2 = \emptyset$ and $y \notin T_2 \supseteq \{x, z\}$. By Lemma A.1, it follows from $Z_i^* \subseteq X_i \cup Y_i$ that

$$x \notin Z_1^* \implies \{y, z\} \subseteq Z_1^* \quad \text{and} \quad y \notin Z_2^* \implies \{x, z\} \subseteq Z_2^*. \quad (4)$$

We now show that neither $x \in Z_1^*$, nor $x \notin Z_1^*$ can be true. This would be the desired contradiction to our assumption that the game possesses a PNE.

For each player $i \in \{1, 2\}$, and each configuration $S_i \in \mathcal{S}_i$, let $c_i^*(S_i)$ denote the price that player i would have to pay so that the resources in S_i are bought, given that the other player $j \in \{1, 2\} \setminus \{i\}$ sticks to her strategy (Z_j^*, p_j^*) . Consider the case $x \notin Z_1^*$: By (4), it follows that $\{y, z\} \subseteq Z_1^*$. Thus, since $Z_1^* \subseteq X_1 \cup Y_1$, the only resources in Z_1^* of non-zero cost are y and z , i.e., $p_1^*(Z_1^*) = p_1^*(y) + p_1^*(z) \leq c_1^*(S_1) = d_1 = 5$. Note that $y \notin Z_2^*$ is not possible, as otherwise player 1 would need to pay $c_1^*(y) = 5\frac{1}{2}$ to buy resource y which is more than the price of $d_1 = 5$ needed to buy S_1 . Thus, $y \in Z_2^*$ must be true. It follows that $p_2^*(z) = 0$, as otherwise $p_2^*(Z_2^*) \geq p_2^*(y) + p_2^*(z) > p_2^*(y) = c_2^*(S_2)$. Thus, since $z \in Z_1^*$, player 1 has to pay $p_2^*(z) = 4$ in order to buy resource z . Since $p_1^*(Z_1^*) \leq c_1^*(S_1) = 5$, it follows that $p_1^*(y) \leq 1$, and therefore, since $y \in Z_2^*$, $p_2^*(y) \geq 4\frac{1}{2}$. However, in this case player 2 could use resource z for free and therefore switch to strategy T_2 for which she would only need to pay the price for resource x which is $d_2 = 4$. So, $x \notin Z_1^*$ is not possible in a PNE.

It remains to consider case $x \in Z_1^*$: Then $p_1^*(Z_1^*) = p_1^*(x) + p_1^*(y) + p_1^*(z) \geq p_1^*(x) = c_1^*(S_1)$ implies $p_1^*(y) = p_1^*(z) = 0$. Therefore $y \notin Z_2^*$, since otherwise, $p_2^*(y) = 5\frac{1}{2}$, so that player 1 could use resource y for free and therefore switch to strategy $T_1 \in \mathcal{S}_1$ of cost $4 = c_z$. However, if $y \notin Z_2^*$, then $\{x, z\} \subseteq Z_2^*$ by equation (4). Hence, $p_2^*(z) = 4$ (since $p_1^*(z) = 0$). It follows that $p_2^*(x) \leq 1\frac{1}{2}$, since otherwise, player 2 would switch to strategy $S_2 \in \mathcal{S}_2$ and pay only the price of $5\frac{1}{2}$ for resource y . Thus, $p_1^*(x) \geq 7\frac{1}{2}$ which is strictly greater than the price of $5\frac{1}{2}$ which player 1 would need to pay if she switches to strategy T_1 . Hence, also $x \in Z_1^*$ is not possible in a PNE, which finishes the proof. \square

A.6 Proof of Theorem 6.1

Proof. Let $p = (p_1, \dots, p_n)$ be the payment vector obtained by marginal cost pricing. Suppose there is a player that unilaterally improves by deviating to some (S'_i, p'_i) . Thus, $S' = (S_1^*, \dots, S_{i-1}^*, S'_i, S_{i+1}^*, \dots, S_n)$, and p'_i differs from p_i only on elements in $S'_i \Delta S_i^* = S'_i \setminus S_i^* \cup$

$S_i^* \setminus S'_i$, while $p'_j = p_j$ for all other players $i \neq j \in N$. For the payoff difference, we therefore calculate that

$$\pi_i(S', p') - \pi_i(S^*, p) = \sum_{r \in S'_i \setminus S_i^*} \left(c_r(\ell_r(S^*) + d_i) - c_r(\ell_r(S^*)) \right) - \sum_{r \in S_i^* \setminus S'_i} p_i^r < 0.$$

Because costs are marginally non-decreasing, we obtain $p_i^r \leq c_r(\ell_r(S^*)) - c_r(\ell_r(S^*) - d_i)$ for all $r \in S_i^*$. Using this inequality we obtain

$$\begin{aligned} c(S') - c(S^*) &= \sum_{r \in S'_i \setminus S_i^*} \left(c_r(\ell_r(S^*) + d_i) - c_r(\ell_r(S^*)) \right) - \sum_{r \in S_i^* \setminus S'_i} \left(c_r(\ell_r(S^*)) - c_r(\ell_r(S^*) - d_i) \right) \\ &\leq \sum_{r \in S'_i \setminus S_i^*} \left(c_r(\ell_r(S^*) + d_i) - c_r(\ell_r(S^*)) \right) - \sum_{r \in S_i^* \setminus S'_i} p_i^r \\ &< 0, \end{aligned}$$

a contradiction to S^* being optimal. □